What is claimed is:

1. An apparatus for routing within a switch fabric, comprising:

a first arbitration component, the first arbitration component configured to receive a first plurality of request-to-sends (RTSs) from a first plurality of fabric gateways including a first fabric gateway, the first arbitration component configured to arbitrate the first plurality of RTSs according to an arbitration scheme; and

a second aribitration component, the second arbitration component configured to receive a second plurality of request-to-sends (RTSs) from a first plurality of fabric gateways including the first fabric gateway, the first arbitration component configured to arbitrate the second plurality of RTSs according to the arbitration scheme.

2. The apparatus of claim 1, wherein:

the first plurality of RTSs are received at the first arbitration component within a time slot; and

the second plurality of RTSs are received at the second arbitration component within a time slot.

3. A method for routing cells within a switch fabric, comprising:

receiving a plurality of cells within a frame, the frame being associated with a plurality of time slots, a subset of cells from the plurality of cells being uniquely associated with each time slot associated with the frame; and

shifting, for each time slot associated with the frame, a frame position for the subset of cells associated with that time slot by an incremental amount from a prior time slot within the frame to produce a shifted frame.

4. The method of claim 3, wherein:

the incremental value is one from an initial value of zero associated with the first time slot,

the frame position for the subset of cells associated with the first time slot being shifted an amount of zero,

the frame position for the subset of cells associated with a first time slot being shifted an amount of one.

5. The method of claim 3, wherein:

the frame being associated with a plurality of rows, each row associated with the frame being associated with an input link,

the shifted frame being associated with a plurality of rows, each row associated with the shifted frame being associated with an output link, each row associated with the shifted frame having a portion of each row associated with the frame.

6. The method of claim 3, wherein:

the shifting is performed within a time delay less than to a total time associated with a time slot from the plurality of time slots for the frame.

7. A method for routing cells within a switch fabric, comprising:

receiving a plurality of cells within a frame, the frame being associated with its own plurality of time slots and its own plurality of rows, each row associated with the frame being associated with an input link; and

reordering the plurality of cells within the frame to produce a shifted frame, the shifted frame having its own plurality of time slots and its own plurality of rows, each row associated with the shifted frame being associated with an output link, each cell being reordered so that each row associated with the frame is uniquely associated with a time slot associated with the shifted frame.

8. The method of claim 7, wherein:

the time slot for each cell within the shifted frame corresponds to the row associated with the frame for that cell.

9. The method of claim 7, wherein:

the row for each cell within the shifted frame corresponds to the time slot associated with the frame for that cell.

10. The apparatus of claim 7, wherein:

the reordering is performed within a time delay less than to a total time associated with the plurality of time slots for the frame.

11. A method for routing cells within a switch fabric, comprising:

time-division multiplexing a plurality of cells associated with a first frame, the first frame being associated with its own plurality of time slots, a plurality of input links and a plurality of output links; and

time-division demultiplexing a plurality of CTSs associated with a second frame, a first CTS from the plurality of CTSs associated with the second frame being associated with an availability of a first RTS associated with a cell from the plurality of cells associated with the first frame.

12. The method of claim 11, wherein the time-division multiplexing includes:

sending each cell from the plurality of cells associated with the first frame to a different switching component from a plurality of switching components; and

receiving each CTS from the plurality of CTSs associated with the second frame from a different switching component from the plurality of switching components.

13. The method of claim 11, further comprising:

time-division multiplexing a plurality of cells associated with a third frame within a delay less than two time slots, the plurality of cells associated with the third frame being associated with its own plurality of time slots, the plurality of input links and the plurality of output links, the plurality of cells associated with the third frame being next in time from the plurality of cells associated with the first frame.

14. An apparatus, comprising:

a plurality of input ports, the plurality of input ports configured to receive a first plurality of cells within a frame, a subset of cells from the plurality of cells being uniquely associated with each time slot associated with the frame;

a cell slot translator coupled to the plurality of input ports, the cell slot translator configured to shift, for each time slot associated with the frame, a frame position for the

subset of cells associated with that time slot by an incremental amount from a prior time slot within the frame to produce a shifted frame; and

a plurality of output ports coupled to the cell slot translator, the plurality of output ports configured to send the shifted frame so that the subset of cells for the first sent time slot of the shifted frame are sent before the subset of cells for the received third time slot of the frame.

15. A method for routing cells within a switch fabric, comprising:

receiving a plurality of cells each having a data portion and a control portion that is unrelated to the data portion for that cell, for each cell from the plurality of cells having a control portion that includes a return-to-send (RTS) that RTS identifying a virtual-output queue (VOQ) having a buffered data portion;

grouping a first plurality of RTSs and a second plurality of RTSs associated with the plurality of received cells to produce a set of grouped RTSs, the first plurality of RTSs being associated with a first frame, the second plurality of RTSs being associated with a second frame different from the first frame; and

arbitrating the set of grouped RTSs to produce a plurality of selected RTSs.

16. The method of claim 15, wherein:

the grouping step produces a vector for each time slot associated with the first frame and the second frame, the vector indicating a status of output port requests for each link from a plurality of links.

17. The method of claim 15, further comprising:

sending a first plurality of clear-to-sends (CTSs) based on the plurality of selected RTSs, each CTS from the first plurality of CTSs being uniquely associated with a corresponding RTS from the plurality of selected RTSs, the first plurality of CTSs including a first CTS that is associated with a first RTS from the plurality of selected RTSs; and

receiving a cell based on a first CTS from the first plurality of CTSs, the received cell having a data portion and a control portion unrelated to the data portion for that received cell,

the data portion of the received cell and the first CTS being associated with the VOQ associated with the first RTS, the received cell being buffered for a time period exceeding other cells buffered within the VOQ associated with the first RTS.

18. The method of claim 15, further comprising:

reducing the set of grouped RTSs based on the plurality of selected RTSs to produce a reduced set of grouped RTSs; and

arbitrating the reduced set of grouped RTSs to produce a second plurality of selected RTSs, the arbitrating of the set of grouped RTSs and the arbitrating of the reduced set of grouped RTSs being performed within a same time slot.

19. The method of claim 15, wherein:

each RTS from the plurality of grouped RTSs is uniquely associated with one output port request from a plurality of output port requests; and

the arbitrating includes:

comparing an output port schedule value with the plurality of output port requests to produce an input port grant;

comparing an input port schedule value with a plurality of input port grants including the produced input port grant to produce an input port/output port designation;

updating the output port schedule value based on the output port of the input port/output port designation; and

updating the input port schedule value based on the input port of the produced input port/output port designation.

20. An apparatus, comprising:

a first switching component, the first switching component being configured to receive a plurality of cells each having a data portion and a control portion that is unrelated to the data portion for that cell, for each cell from the plurality of cells having a control portion that includes a return-to-send (RTS) that RTS identifying a virtual-output queue (VOQ) having a buffered data portion,

the first switching component being further configured to group a first plurality of RTSs and a second plurality of RTSs associated with the plurality of received cells to

produce a set of grouped RTSs, the first plurality of RTSs being associated with a first frame, the second plurality of RTSs being associated with a second frame different from the first frame,

the first switching component being further configured to arbitrate the set of grouped RTSs to produce a plurality of selected RTSs.

21. The apparatus of claim 20, wherein:

the first switching component is configured to reduce the set of grouped RTSs based on the plurality of selected RTSs to produce a reduced set of grouped RTSs, the first switch is configured to arbitrate the reduced set of grouped RTSs to produce a second plurality of selected RTS.

22. The apparatus of claim 20, wherein:

each RTS from the plurality of grouped RTSs is uniquely associated with one output port request from a plurality of output port requests;

the first switching element compares an output port schedule value with the plurality of output port requests to produce an input port grant;

the first switching element compares an input port schedule value with a plurality of input port grants including the produced input port grant to produce an input port/output port designation;

the first switching element updates the output port schedule value based on the output port of the input port/output port designation; and

the first switching element updates the input port schedule value based on the input port of the produced input port/output port designation.

23. An apparatus associated with a plurality of links, comprising:

a first memory component, the first memory component configured to group a plurality of output port requests;

a second memory component coupled to the first memory component, the second memory component being configured to store a vector for each time slot associated with a frame, the vector for each time slot indicating a status of output port requests for each link from the plurality of links;

a plurality of comparators coupled to the second memory component, each comparator from the plurality of comparators being associated with an input port and an output port, each comparator from the plurality of comparators being configured to compare an input port schedule value with the plurality of input port requests to produce an output port grant, each comparator from the plurality of comparators being further configured to compare an output port schedule value with a plurality of output port grants including the produced output port grant to produce an input port/output port designation; and

a third memory component coupled to the plurality of comparators, the third memory component being configured to store the input port/output port designation produced by each comparator from the plurality of comparators.

24. The apparatus of claim 23, wherein:

each vector stored within the second memory is associated with an input port designation and indicates a validity/priority value for each combination of the input portion designation and an output port designation,

the input port designation and the output port designation being associated with a given link from the plurality of links.

25. The apparatus of claim 23, further comprising:

a parser coupled to the plurality of comparators, the parser configured to receive a plurality of cells each having a data portion and a control portion that is unrelated to the data portion, for each cell from the plurality of cells having a control portion that includes an RTS, that RTS indicating a virtual-output queue (VOQ) having a buffered data portion, the parser further configured to separate the data portion associated with each cell from the plurality of cells; and

a fourth memory component coupled to the parser, the fourth memory component configured to store the data portions associated with the cells from the plurality of cells.

26. A method for routing cells within a switch fabric, comprising:

grouping a plurality of request-to-sends (RTSs);

forming a plurality of vectors based on the grouped RTSs, each vector from the plurality of vectors being associated with a time slot within a frame, the vector for each time slot indicating a status of output port request for each link from a plurality of links; and

arbitrating the plurality of RTSs based on the plurality of vectors to produce a set of input port/output port designations.

27. The method of claim 26, wherein each cell from a plurality of cells includes a data portion and a control portion unrelated to the data portion for that cell, for each cell from the plurality of cells having a control portion that includes an RTS from the plurality of RTSs , the arbitrating being performed with that RTS without the associated data portion.

28. An apparatus, comprising:

a first switching component having a grouping memory, an arbitration component and its own plurality of input ports and its own plurality of output ports, the first switching component receiving a first plurality of request-to-sends (RTSs) at its plurality of input ports within a first frame having its own plurality of time slots, the first plurality of RTSs being stored in the grouping memory, the arbitration component arbitrating concurrently the first plurality of RTSs to produce a plurality of selected RTSs.

29. The apparatus of claim 28, wherein:

each cell from a plurality of cells includes a data portion and a control portion unrelated to the data portion for that cell, for each cell from the plurality of cells having a control portion that includes an RTS from the plurality of RTSs, the arbitration component arbitrating the first plurality of RTSs without the data portion associated with each RTS from the first plurality of RTSs.

30. The apparatus of claim 28, wherein:

the grouping memory configured to reduce the first pluarlity of RTSs based on the plurality of selected RTSs,

the first switching component configured to receive a second plurality of request-to-sends (RTSs) at its plurality of input ports within a second frame having its own plurality of time slots, the second frame being different from the first frame,

the arbitration component arbitrating concurrently the reduced first plurality of RTSs and the second plurality of RTSs to produce a second plurality of selected RTSs.

31. The apparatus of claim 28, further comprising:

a second switching component;

a first switch fabric gateway having its own plurality of input ports and its own plurality of output ports, a first output port from the plurality of output ports associated with the first switch fabric gateway being coupled to the first switching component, a second output port from the plurality of output ports associated with the first switch fabric gateway being coupled to the second switching component; and

a second switch fabric gateway having its own plurality of input ports and its own plurality of output ports, a first output port from the plurality of output ports associated with the second switch fabric gateway being coupled to the switching component, a second output port from the plurality of output ports associated with the second switch fabric gateway being coupled to the second switching component

32. The apparatus of claim 28, further comprising:

a second switching component;

a first switch fabric gateway having its own plurality of input ports and its own plurality of output ports, a first input port from the plurality of input ports associated with the first switch fabric gateway being coupled to the first switching component, a second input port from the plurality of input ports associated with the first switch fabric gateway being coupled to the second switching component; and

a second switch fabric gateway having its own plurality of input ports and its own plurality of output ports, a first input port from the plurality of input ports associated with the second switch fabric gateway being coupled to the switching component, a second

input port from the plurality of input ports associated with the second switch fabric gateway being coupled to the second switching component.

33. An apparatus coupled to a plurality of cell framers including a first cell framer and a second cell framer, comprising:

a buffer memory;

a data storage controller coupled to the buffer memory, the data storage controller being configured to receive:

a first data cell associated with a first time slot from a first cell framer, and

a second data cell associated with the first time slot from a second cell framer; and

a data alignment controller coupled to the data storage controller, the data alignment controller being configured to send a forwarding signal to the data storage controller after a data cell for the first time slot has been received from each cell framer that is active from the plurality of cell framers.

34. The apparatus of claim 33, wherein:

the first data cell has an associated link-status marker, and

the data storage controller is further configured to replace a value associated with the first data cell with an idle value when the link-status marker associated with the first data cell has a bad-link value.

35. The apparatus of claim 33, wherein:

the first data cell has an associated cell-status marker, and

the data storage controller is further configured to replace a value associated with the first data cell with an idle value when the cell-status marker associated with the first data cell has a bad-cell value.

36. The apparatus of claim 33, wherein the data storage controller is further configured to:

receive a plurality of data cells including the first data cell and the second data cell, each cell from the plurality of data cells having an associated cell framer identifier from a plurality of cell framer identifiers associated with the data storage controller, and

postpone forwarding data cells from the plurality of data cells upon receiving a data cell having an associated start-of-super-frame value until a data cell having an associated start-of-super-frame value is received for each cell framer identifier from the plurality of cell framers identifiers.

37. An apparatus within a switching fabric, comprising:

a buffer memory;

a data storage controller coupled to the buffer memory, the data storage controller being configured to receive a plurality of data cells associated with a first time slot, each data cell from the plurality of data cells being uniquely associated with its own cell framer and its own receipt time; and

a data alignment controller coupled to the data storage controller, the data alignment controller being configured to send a forwarding signal to the data storage controller at the latest receipt time associated with the plurality of data cells that is within a timeout period.

38. The apparatus of claim 37, wherein:

the plurality of data cells includes a first data cell, the first data cell has an associated link-status marker, and

the data storage controller is further configured to replace a value associated with the first data cell with an idle value when the link-status marker associated with the first data cell has a bad-link value.

39. The apparatus of claim 37, wherein the data storage controller is further configured to:

postpone forwarding data cells from the plurality of data cells associated with the first time slot upon receiving a data cell having an associated start-of-super-frame value until a data cell having an associated start-of-super-frame value is received for each cell framer identifier from the plurality of cell framers identifiers.

40. The apparatus of claim 37, wherein the data controller is configured to provide an idle cell for each cell from the plurality of cells that is not received within the timeout period.

41. An apparatus, comprising:

receiving a plurality of cells associated with a first time slot, each cell from the plurality of cells being uniquely associated with its own cell framer and its own receipt time;

buffering each cell from the plurality of cells as that cell is received; and

sending the plurality of cells after every cell from the plurality of cells are received within a timeout period.

42. The apparatus of claim 41, further comprising:

providing, before sending the plurality of cells, an idle cell for each cell from the plurality of cells that are not received within the timeout period.

43. A switching fabric, comprising:

a first physical switch stage having its own plurality of components; and

a second physical switch stage having its own plurality of components, the plurality of components associated with the second physical switch stage being coupled to the plurality of components associated with the first physical switch stage,

the plurality of components associated with the second physical switch stage including at least a first component and a second component, the first component and the second component each having its own deskew module configured to forward substantially simultaneously a first data cell associated with a first time slot and a first receipt time, and a second data cell associated with the first time slot and a second receipt time later than the first receipt time.

44. The switching fabric of claim 43, wherein:

the first component associated with the first physical switch stage has its own associated clock speed; and

the second component associated with first physical switch stage has its own associated clock speed different from the clock speed associated with the first component.

45. The switching fabric of claim 43, further comprising:

a third physical switch stage having its own plurality of components including a first component and a second component, the plurality of components associated with the third physical switch stage being coupled to the plurality of components associated with the second physical switch stage,

the first component associated with the first physical switch stage has its own associated clock speed,

the first component associated with second physical switch stage has its own associated clock speed different from the clock speed associated with the first component associated with the first physical switch stage.

46. The switching fabric of claim 43, further comprising:

a plurality of connections including a first connection and a second connection, the plurality of components associated with the second physical switch stage being coupled to the plurality of components associated with the first physical switch stage by the plurality of connections,

the first connection having a length and coupling the first component of the first physical switch stage and the first component of the second physical switch stage, the second connection having a length and coupling the second component of the first physical switch stage to the first component of the second physical switch stage, the length of the first connection being different from the length of the second connection.

47. The switching fabric of claim 43, wherein:

the first data cell has an associated link-status marker, and

the first component associated with the first physical switch stage is configured to replace a value associated with the first data cell with an idle value when the link-status marker associated with the first data cell has a bad-link value.

48. The switching fabric of claim 43, wherein:

the first data cell has an associated cell-status marker, and

the first component associated with the first physical switch stage is configured to replace a value associated with the first data cell with an idle value when the cell-status marker associated with the first data cell has a bad-cell value.

49. The switching fabric of claim 43, wherein:

the first data cell has an associated start-of-super-frame marker, and

the first component associated with the first physical switch stage is configured to replace a value associated with the first data cell with an idle value when the link-status marker associated with the first data cell has a bad-link value.

50. The switching fabric of claim 43, wherein:

the deskew module for the first component is configured to postpone forwarding any cells upon receiving a cell having an associated start-of-super-frame value until a cell having an associated start-of-super-frame value is received for each component from the first physical switch stage coupled to the first component of the second physical switch stage.

51. A switching fabric, comprising:

a first plurality of line cards;

a first plurality of multiplexer/demultiplexer cards coupled to the first plurality of line cards; and

a first plurality of switching cards removably coupled to and geometrically reconfigurable with the first plurality of multiplexer/demultiplexer cards.

52. The switching fabric of claim 51, further comprising:

a first plurality of optical fibers, the first plurality of switching cards being removably coupled to and geometrically reconfigurable with the first plurality of multiplexer/demultiplexer cards by the first plurality of optical fibers.

53. The switching fabric of claim 51, further comprising:

a second plurality of line cards;

a second plurality of multiplexer/demultiplexer cards coupled to the first plurality of line cards and the second plurality of line cards, the second plurality of line cards being coupled to the first plurality of multiplexer/demultiplexer cards; and

a second plurality of switching cards removably coupled to and geometrically reconfigurable with the first plurality of multiplexer/demultiplexer cards and the second plurality of multiplexer/demultiplexer cards.

54. The switching fabric of claim 51, further comprising:

a second plurality of optical fibers, the first plurality of switching cards being removably coupled to and geometrically reconfigurable with the first plurality of multiplexer/demultiplexer cards and the second plurality of multiplexer/demultiplexer cards by the first plurality of optical fibers, the second plurality of switching cards being removably coupled to and geometrically reconfigurable with the first plurality of multiplexer/demultiplexer cards and the second plurality of multiplexer/demultiplexer cards by the second plurality of optical fibers.

55. A switching fabric, comprising:

a first plurality of fabric gateway components;

a first plurality of multiplexer/demultiplexer components coupled to the first plurality of fabric gateway components;

a second plurality of fabric gateway components;

a second plurality of multiplexer/demultiplexer components coupled to the second plurality of fabric gateway components; and

a plurality of switching components removably coupled to the first set of multiplexer/demultiplexer components and the second set of multiplexer/demultiplexer components.

56. The switching fabric of claim 55, wherein:

the first plurality of multiplexer/demultiplexer components include at least a first portion and a second portion;

the second plurality of multiplexer/demultiplexer components include at least a first portion and a second portion; and

the plurality of switching components include at least a first portion and a second portion, the first portion of the first plurality of multiplexer/demultiplexer components and the first portion of the second plurality of multiplexer/demultiplexer components being removably coupled to the first portion of the plurality of switching components, the second portion of the first plurality of multiplexer/demultiplexer components and the second portion of the second plurality of multiplexer/demultiplexer components being removably coupled to the second portion of the plurality of switching components.

57. The switching fabric of claim 55, wherein:

a plurality of switching components are removably coupled to the first plurality of multiplexer/demultiplexer components and the second plurality of multiplexer/demultiplexer components by a plurality of optical fibers.

58. The switching fabric of claim 55, wherein:

a plurality of switching components are geometrically reconfigurable with the first plurality of multiplexer/demultiplexer components and the second plurality of multiplexer/demultiplexer components by a plurality of optical fibers.

59. The switching fabric of claim 55, wherein:

a plurality of switching components are removably coupled to the first plurality of multiplexer/demultiplexer components and the second plurality of multiplexer/demultiplexer components regardless of a geometric configuration of the plurality of switching components, the first plurality of multiplexer/demultiplexer components and the second plurality of multiplexer/demultiplexer components.

60. The switching fabric of claim 55, wherein:

a first switching component from the plurality of switching components has its own plurality of ports;

at least a first port from the plurality of ports for the first switching component being removably coupled to a first multiplexer/demultiplxer component from the plurality of multiplexer/demultiplxer components by a first optical fiber; and

at least a second port from the plurality of ports for the first switching component being removably coupled to a second multiplexer/demultiplxer component from the plurality of multiplexer/demultiplxer components by a second optical fiber.

61. The switching fabric of claim 55, wherein:

a first switching component from the plurality of switching components has its own plurality of ports,

each port from the plurality of ports for the first switching component being removably coupled to a different multiplexer/demultiplxer component from the plurality of multiplexer/demultiplxer components by a different optical fiber.

62. A switching fabric, comprising:

a plurality of fabric gateway components; and

a plurality of configurable components each having a first configuration and a second configuration, the first configuration being a multiplexer/demultiplexer, the second configuration being a switching component,

a first set of configurable components from the plurality of configurable components each being configured as the first configuration, the first set of configurable components being coupled to the plurality of fabric gateway components,

a second set of configurable components from the plurality of configurable components each being configured as the second configuration, the second set of configurable components being removably coupled to the plurality of the first set of configurable components.

63. The switching fabric of claim 62, wherein:

the second set of configurable components are removably coupled to the first set of configurable components by a plurality of optical fibers.

64. The switching fabric of claim 62, wherein:

the second set of configurable components are removably coupled in geometrically reconfigurable manner to the first set of configurable components.

65. A method for configuring a switching fabric having at least a first plurality of configurable components coupled to its own plurality of line cards and having a first configuration and a second configuration, the first configuration being a multiplexer/demultiplexer, the second configuration being a switching component, comprising:

providing a second plurality of configurable components being in the second configuration;

reconfiguring the first plurality of configurable components from the second configuration to the first configuration; and

removably coupling the second plurality of configurable components to the first plurality of configurable components.

66. The method of claim 65, wherein:

the second set of configurable components are removably coupled to the first set of configurable components by a first plurality of optical fibers.

67. The method of claim 65, further comprising:

providing a third plurality of configurable components being in the first configuration and being coupled to its own plurality of line cards;

uncoupling the first plurality of configurable components from the second plurality of configurable components; and

removably coupling the first plurality of configurable components to the second plurality of configurable components and the third plurality of configurable components.

68. The method of claim 65, further comprising:

providing a third plurality of configurable components being in the first configuration and being coupled to its own plurality of line cards;

uncoupling the first plurality of optical fibers from the second plurality of configurable components;

providing a second plurality of optical fibers; and

removably coupling the first plurality of configurable components to the second plurality of configurable components and the third plurality of configurable components by the first plurality of optical fibers and the second plurality of optical fibers.

69. A method for expanding a switching fabric having at least a first plurality of multiplexer/demultiplexer cards coupled to its own plurality of line cards, and a second plurality of multiplexer/demultiplexer cards coupled to its own plurality of line cards, the first plurality of multiplexer/demultiplexer cards and the second plurality of multiplexer/demultiplexer cards being removably coupled to a first plurality of switching cards, the first plurality of switching cards each having at least one switching component, said method comprising:

uncoupling the first plurality of switching cards from the first plurality of multiplexer/demultiplexer cards and the second plurality of multiplexer/demultiplexer cards;

providing at least a third plurality of multiplexer/demultiplexer cards coupled to its own plurality of line cards;

providing a second plurality of switching cards; and

removably coupling the first plurality of switching cards and the second plurality of switching cards to the first plurality of multiplexer/demultiplexer cards, the second plurality of multiplexer/demultiplexer cards and the third plurality of multiplexer/demultiplexer cards.

70. The method of claim 69, further comprising:

a plurality of optical fibers, the first plurality of switching cards and the second plurality of switching cards being removably coupled by the plurality of optical fibers to

the first plurality of multiplexer/demultiplexer cards, the second plurality of multiplexer/demultiplexer cards and the third plurality of multiplexer/demultiplexer cards.

71. A method for expanding a switching fabric having at least a first plurality of multiplexer/demultiplexer component coupled to its own plurality of fabric gateway components, and a second plurality of multiplexer/demultiplexer components coupled to its own plurality of fabric gateway components, the first plurality of multiplexer/demultiplexer components and the second plurality of multiplexer/demultiplexer components being removably coupled to a first plurality of switching components, said method comprising:

uncoupling the first plurality of switching components from the first plurality of multiplexer/demultiplexer cards and the second plurality of multiplexer/demultiplexer cards;

providing at least a third plurality of multiplexer/demultiplexer components coupled to its own plurality of fabric gateway components;

providing a second plurality of switching components;

removably coupling the first plurality of switching components and the second plurality of switching components to the first plurality of multiplexer/demultiplexer components, the second plurality of multiplexer/demultiplexer components and the third plurality of multiplexer/demultiplexer components.

72. The method of claim 71, wherein:

the first plurality of switching components and the second plurality of switching components are removably coupled to the first plurality of multiplexer/demultiplexer components, the second plurality of multiplexer/demultiplexer components and the third plurality of multiplexer/demultiplexer components by a plurality of optical fibers.

73. A method for randomizing a plurality of request-to-sends (RTSs) within a switching fabric, each RTS from the plurality of RTSs being associated with a destination identifier from a plurality of destination identifiers, said method comprising:

performing, for each time slot associated with a frame, the following:

randomly selecting a link from a plurality of links for that time slot;

inserting a first RTS from the plurality of RTSs into a cell associated with the selected link for that time slot; and

inserting, for each link from the plurality of links for that time slot, a remaining RTSs from the plurality of RTSs into a respective cell associated with that time slot.

74. An apparatus, comprising:

a plurality of input ports;

a plurality of output ports; and

a switching fabric coupled to the plurality of input ports and the plurality of output ports, the switching fabric having a distributed scheduler configured to schedule the routing of a plurality of cells from the plurality of input ports to the plurality of output ports, the distributed scheduler having a control path with its own rate and a data path with its own rate, the rate of the control path of the distributed scheduler being less than a rate of a control path of a centralized scheduler with a data path having a rate similar to the data rate of the distributed scheduler.

75. The apparatus of claim 74, wherein:

each cell from the plurality of cells having a data portion and a control portion that is unrelated to the data portion for that cell, the control portion of each cell from the plurality of cells being associated with the control path of the distributed scheduler, the data portion of each cell from the plurality of cells being associated with the data path of the distributed scheduler.

76. A method for routing cells within a switching fabric, comprising:

receiving a plurality of cells each being associated with a priority value from a plurality of priority values, the plurality of priority values including a first priority value and a second priority value;

buffering the plurality of cells in a plurality of virtual output queues (VOQs) including a first VOQ, the first VOQ being associated with the first priority value and the second priority value, each remaining VOQ from the plurality of VOQs being uniquely associated with a remaining priority value from the plurality of priority values, each cell

from the plurality of cells being buffered in the VOQ that is associated with the priority value matching the prior value of that cell; and

sending, for each cell buffered in the first VOQ, a new-cell indication signal having the first priority value when a queue length of the first VOQ is above a predefined threshold and having the second priority value when the queue length of the first VOQ is not above the predefined threshold.

77. The method of claim 76, wherein:

the first priority value is associated with a minimum bit rate service, and

the second priority value is associated with a value reduced from the first priority value.

78. The method of claim 76, further comprising:

sending a flow-control signal associated with the second priority when the queue length of the first VOQ is above the predefined threshold; and

receiving, based on the sent flow-control signal, cells associated with the first priority value at a rate not greater than a minimum bit rate service.

79. An apparatus, comprising:

a plurality of virtual output queues (VOQs) including a first VOQ, the first VOQ being associated with a first priority value and a second priority value from a plurality of priority values, each remaining VOQ from the plurality of VOQs being uniquely associated with a remaining priority value from a plurality of priority values;

a first counter associated with the first VOQ and the first priority value, the first counter being configured to count each cell being buffered in the first VOQ and having the first priority value;

a second counter associated with the first VOQ and the second priority value, the second counter being configured to count each cell being buffered in the first VOQ and having the second priority value; and

a signal generator coupled to the first VOQ, the signal generator being configured to send, for each cell buffered in the first VOQ, a new-cell indication signal having the first priority value when a count of the first counter is above a predefined threshold and

having the second priority value when a count of the second counter is not above the predefined threshold.

80. The apparatus of claim 79, wherein:

the first priority value is associated with a minimum bit rate service, and

the second priority value is associated with a value reduced from the first priority value.

81. The apparatus of claim 79, wherein:

the first counter and the second counter are disposed within an RTS generator;

the plurality of VOQs is disposed within a VOQ manager; and

the signal generator is disposed within a flow controller, the flow controller being coupled to the RTS generator and the VOQ manager.

82. The apparatus of claim 79, wherein:

the signal generator being configured to send a flow-control signal associated with the second priority when the queue length of the first VOQ is above the predefined threshold; and

the pluarlity of VOQs receiving, based on the sent flow-control signal, cells associated with the first priority value at a rate not greater than a minimum bit rate service.